

This is an extended version of an article originally published online in Data Storage Magazine, January 2002.

# Deep Memory Scopes Speed BER Analysis

By Charles H. Sobey, ChannelScience  
and Jay Alexander, Agilent Technologies

The ultimate performance measure in most data communication and data storage systems is the *bit error rate* (BER). BER is defined as the number of bits detected in error divided by the total number of bits detected. Typically, this is averaged over many different transmitted (or stored) data sequences. As such, BER is a measure of the *probability of error*. Speed, power and cost are also important system-level considerations. However, if the information that is being sent or stored at high speed, low power and low cost is not correct, these other metrics are meaningless.

Exactly what value of BER is considered to be “good,” that is low, depends on the application. For example, a crude voice application, in which frequent re-send requests are acceptable, may be able to operate at a BER of  $10^{-3}$  or  $10^{-4}$ . Whereas a hard disk drive typically has a BER around  $10^{-14}$  to  $10^{-16}$ .

How can such low BERs be tested? 50 to 200 bit errors are typically sought before a statistically significant BER can be measured. Approximately  $10^{16}$  bits would have to be detected to measure a  $10^{-14}$  BER. Even at the highest disk drive data rates (about 1 gigabit per second) this would take  $10^7$  seconds, which is over 1157 days. The test probably lasts longer than the warranty!

## Raw BER and the Role of ECC

The very low BER values used in the example above are typically available only *after* applying aggressive error correction coding (ECC). This is a deterministic calculation, which exploits structured redundancy that was encoded into the data when it was transmitted or stored. ECC can provide several orders of magnitude improvement (reduction) in BER. Therefore, for speed, it is common to test and measure BER *before* the ECC's corrections are applied.<sup>1</sup> This is often referred to as the *raw* BER.

In the disk drive example above the raw BER may be around  $10^{-9}$ . To observe about 100 bit errors, at a data rate of 1 gigabit per second, the test time would be about 100 seconds. Such testing may be *possible*, but it is not at all economically practical. Depending on the application, a BER test might be allotted a few seconds or as little as a fraction of a second. In addition, it may be necessary to test BER under several sets of conditions. For communication systems these may include different numbers of users, signal strengths, data rates, fading characteristics and distortions. Because of these considerations, it is still desirable to reduce the BER test time further.

## The Fastest BER Test is NO Test at All

Of course, the best way to minimize BER test time is to eliminate the test. However, this can only be accomplished by ensuring that there is adequate performance margin built into the system. This requires that the system and all of its components be properly designed for system-level (BER) performance -- preferably from the very beginning of the project.

---

<sup>1</sup> Care must be taken with this method because often the *distribution* of errors affects ECC performance.

To speed BER testing, it is common to stress the system to increase the probability that bit errors will occur. This is usually done by testing under “worst-case” conditions or by “de-tuning” key system parameters. Worst-case conditions include extremes of noise, multiple users, cross talk, echo, fading, or, in the case of hard disk drives, positioning of the read element off-track relative to the written magnetic track. Calibrating and correlating worst-case conditions across multiple testers is difficult. Candidates for parameter de-tuning include key signal processing steps such as equalization, non-linearity compensation, phase-locked loop timing control or detection parameters. This requires the component suppliers to have intimate knowledge of the signal processing hardware, which they otherwise would not need.

These methods test the system and the components under rather artificial extremes. This can lead to inappropriate design choices that degrade system performance under more typical, nominal conditions. How can component suppliers quickly and accurately test for BER, under realistic conditions, without having intimate knowledge of the system's signal processing?

## Deep Memory Scopes Break the BER Testing Bottleneck for Components

It is possible to design components based on system-level performance without the use of exact hardware channels, without excessive calibration and correlation issues and without resorting to unrealistic worst-case scenarios. The solution is a combination of a signal processing/detection simulation, BER projection and a deep memory digital sampling oscilloscope.

Any component that is crucial to the generation of an analog or digital signal that eventually must be processed and detected by the system is a candidate for this type of testing. This could include the heads, media and preamplifier of data storage devices or the laser, lenses and receiver of a fiber optic link. It could also be the output of a signal conditioning block, such as a regenerator or a filter. Because of bandwidth considerations, the best results are obtained with baseband signals.

The first step is to capture the analog signal of the component under test with a fast, deep memory scope. As a rule-of-thumb, the sampling rate of the scope should be at least 4 times the data rate of the signal of interest. In practice, a factor of 6 to 10 times is more than sufficient. So a 4 Gsample/second oscilloscope is perhaps not adequate for a 1Gbps data rate signal, but an 8 Gsample/second scope is.

To measure a BER of  $10^{-8}$ , over 40 billion samples will be required. Clearly deep memory is essential in order to capture a large number of bits at one time. With typical digital sampling oscilloscopes, one must tradeoff sampling rate for memory depth. Figure 1 shows a typical readback signal captured from a hard disk drive preamplifier. Agilent's latest Infiniium [1] scope was used to capture the waveform because its “MegaZoom®” feature provides full memory depth at even the highest sampling rates. However, *no* scope is within even 3 orders-of-magnitude of having 40 billion points of fast memory! So how can accurate estimates of low BERs be obtained?

## Scope + Simulation Provide what is NOT Available in Hardware!

Under the right conditions, it is possible to *project* the BER to low values using orders-of-magnitude less samples than required for a direct measurement. The key is to have access to the detector's internal decision metric. However, this is seldom available from hardware. For this reason, a software simulation is needed. The simulation must adequately reflect the signal processing chain of the hardware channel and must employ a very similar detection algorithm.

Consider the well-known Viterbi detector, as it is applied in hard disk drives and some tape and optical drives. In particular, it is a partial response maximum likelihood (PRML) sequence detection system [2]. The Viterbi algorithm uses an internal decision variable known as the *difference metric* [3]. For the purposes of this article, the difference metric is defined as the absolute value of the difference in squared error between the two sample sequences (“paths”) that the Viterbi detector decides between, at each bit along the chosen path. The upper-left plot in Figure 4 shows an example scatter plot of the Viterbi

difference metric. The lower-left plot is the corresponding distribution (histogram), which is used for BER projection. The blue S-shaped curve is the cumulative sum of the histogram, normalized to a maximum value of 1.

## Projecting BER

If the difference metric distribution is gaussian, the normalized cumulative sum approaches the shape of the well-known gaussian *error function*,  $\text{erf}(x)$ , as the number of samples increases. The error function directly relates to the cumulative probability under a gaussian probability density function from  $x = -\infty$  to  $x = +\infty$ .

Through the probability, each value of  $x$  maps into a corresponding number of standard deviations,  $\sigma$ . If the difference metric distribution is truly gaussian, the relationship between  $x$  and the number of  $\sigma$ s from the mean is linear, as shown in the lower-right plot of Figure 4. Colored noise can cause the difference metric distribution to be non-gaussian, resulting in a poor “linear” fit (low  $r^2$ ).

Each triangle in the plot represents a histogram bin ( $x$ -axis of distribution plot) that contains 10 or more occurrences of difference metric values. If a line is fit to these triangles and extended to intersect the  $y$ -axis ( $x = 0$ ), the  $y$ -intercept projects the ultimate performance of the system. This is in terms of how many  $\sigma$ s of the difference metric distribution can be detected error free. Only the histogram bins that contain  $\geq 10$  difference metric occurrences are used to make the linear fit. The slope of the fit reflects the value of  $\sigma$ . The upper-right plot of Figure 4 shows the calculated BER corresponding to the number of  $\sigma$ s and the linear fit.

## Example: Projecting BER Using a Captured Hard Disk Drive Signal

Figure 2a shows the same hard disk drive readback signal as Figure 1, plus its histogram. The Infinium's orange markers, Ax and Bx, delimit the portion of the waveform that is included in the histogram. The waveform out of the read channel's continuous-time filter is displayed in Figure 2b, along with its histogram. The goal of the continuous-time (and discrete-time) filtering is to equalize this waveform to five distinct target amplitude levels, at synchronously sampled instants.

The five more-distinct peaks in the distribution in Figure 2b are matched much better to the five target sample values expected by the Viterbi detector. In the hardware channel, this filtered waveform will be sampled and then filtered again by the FIR of the read channel before being passed to the detector. If the distribution of sample values at the input to the detector were available it would show five even more pronounced peaks.

To obtain the difference metric distribution, the continuous-time and discrete-time signal processing steps are modeled using ChannelScience's PRML*pro*<sup>TM</sup> read channel simulation package [4]. The channel front-end screen shown in Figure 3 provides integrated models of the most common signal processing steps performed in data storage devices. These steps include ac coupling, 3<sup>rd</sup> order non-linearity compensation, automatic gain control (AGC), continuous-time filtering, phase and frequency locked loop (PLL) control of sampling, discrete-time filtering using an adaptive finite impulse response (FIR) filter and whitener.

The resulting samples are then processed by PRML*pro*'s Viterbi detector simulation. The difference metrics are saved for BER projection as described above. The lower-right plot of Figure 4 shows that the  $y$ -intercept ( $x = 0$ ) of the linear fit is 5.353. The upper-right plot of Figure 4 shows the corresponding BER for this projected number of  $\sigma$ s is  $10^{-7.36}$ . To measure this BER would require detecting about 40 billion bits! Only 81,531 were used for this projection. However, it is important to be careful with sample size.

Notice that the triangles in this plot stop at a BER of about  $10^{-3.23}$ . To get to  $10^{-7.36}$  we must project 4.13 orders of magnitude. The fewer orders of magnitude that the projection spans, the more confidence we can have in it. In general, projections that span two to three orders of magnitude should be reliable.

# ChannelScience.com

---

Beyond three and half or four orders, it is advisable to take a larger sample size. In our example, a longer waveform should be captured and simulated in order to have more confidence in the projection. Again, this is where the deep memory and the high sampling rate of the scope are crucial.

Agilent's Infiniium oscilloscope has an internal PC that runs MS Windows. Therefore PRML<sup>pro</sup>™ can be run right on the scope's display making for a compact, convenient test combination where lab bench space is tight. Once the simulation and projection are qualified for a particular application, the same software is used on every scope, which reduces calibration and correlation issues to simply following good scope maintenance practices.

## References

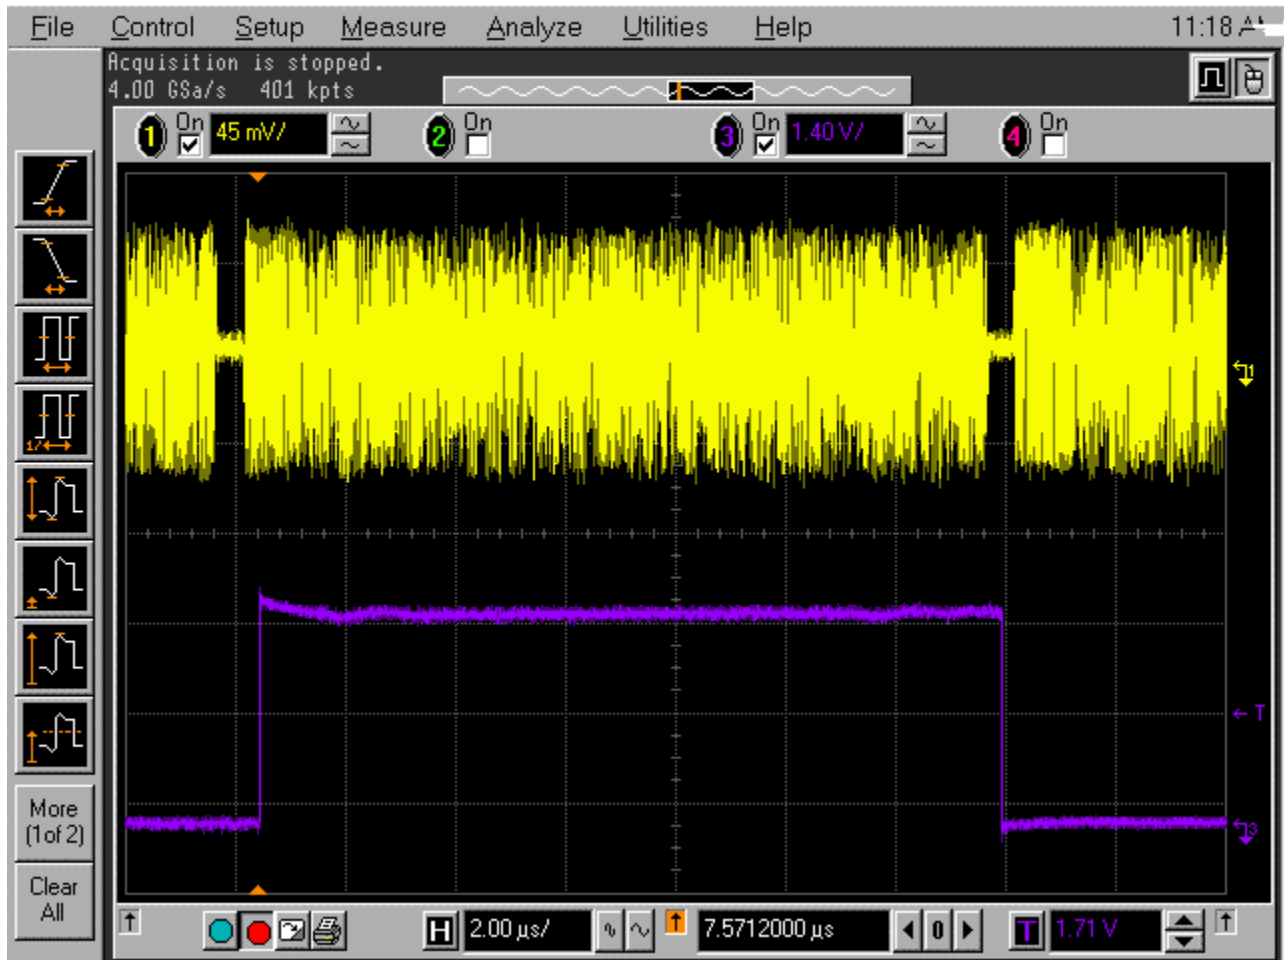
- [1] Agilent Technologies 54832B oscilloscope data sheet, [www.agilent.com/find/infiniium](http://www.agilent.com/find/infiniium).
- [2] R.D.Cideciyan, *et al.*, "A PRML System for Digital Magnetic Recording," *IEEE Journal on Selected Areas in Communications*, Vol. 10, No. 1, January 1992, pp. 38-56.
- [3] T.Perkins and Z. Keirn, "A Window-margin-like Procedure for Evaluating PRML Channel Performance," *IEEE Transactions on Magnetics*, Vol. 31, No. 2, March 1995, pp. 1109-1114.
- [4] Details about PRML<sup>pro</sup>™, including a downloadable user manual and free trial software, are available at [www.ChannelScience.com](http://www.ChannelScience.com).

## Author Bios

Chuck Sobey ([csobey@channelscience.com](mailto:csobey@channelscience.com)) is an internationally known technical consultant, lecturer and business advisor. As Chief Scientist of ChannelScience.com he directs and contributes to the development of software, signal processing algorithms and chip architectures for data storage, data communications and bioinstrumentation devices. Recent research activities include the areas of turbo coding, medical signal processing and cochannel detection for sensor arrays.

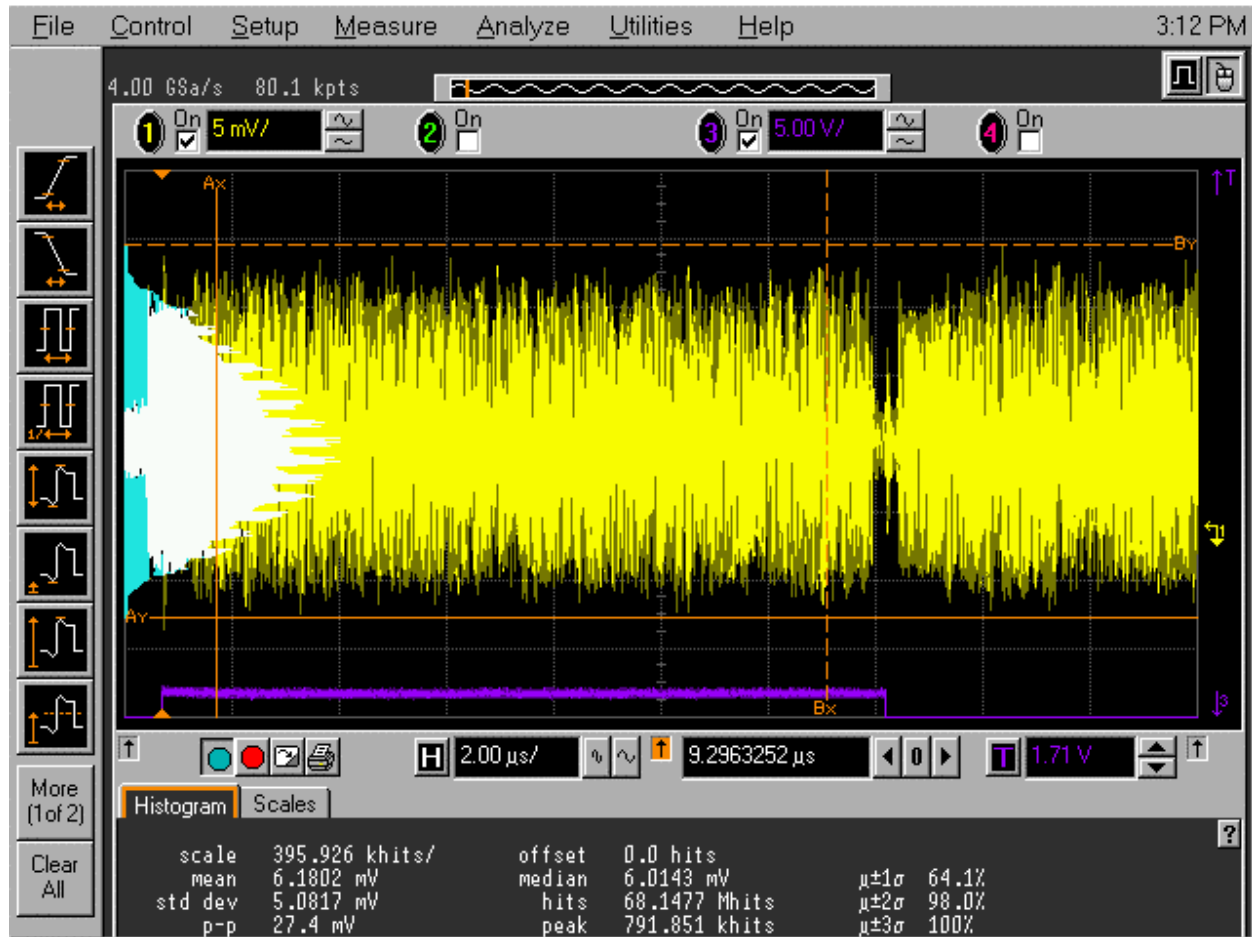
Jay Alexander is an R&D Section Manager responsible for oscilloscope development at Agilent Technologies. He is a licensed Professional Engineer in Colorado and is a Senior Member of IEEE. He designed the HP/Agilent Infiniium oscilloscope user interface and holds several patents for that work.

Saved: 30 OCT 2001 11:18:54




**Figure 1.** The yellow waveform is the hard disk drive readback signal at the output of the preamplifier. The purple waveform is the “Read Gate.” This signal indicates to the drives read channel when to begin and end reading the input waveform. This is one “sector” of the disk drive, which typically stores 512 bytes of user data.

Saved: 30 OCT 2001 15:12:26

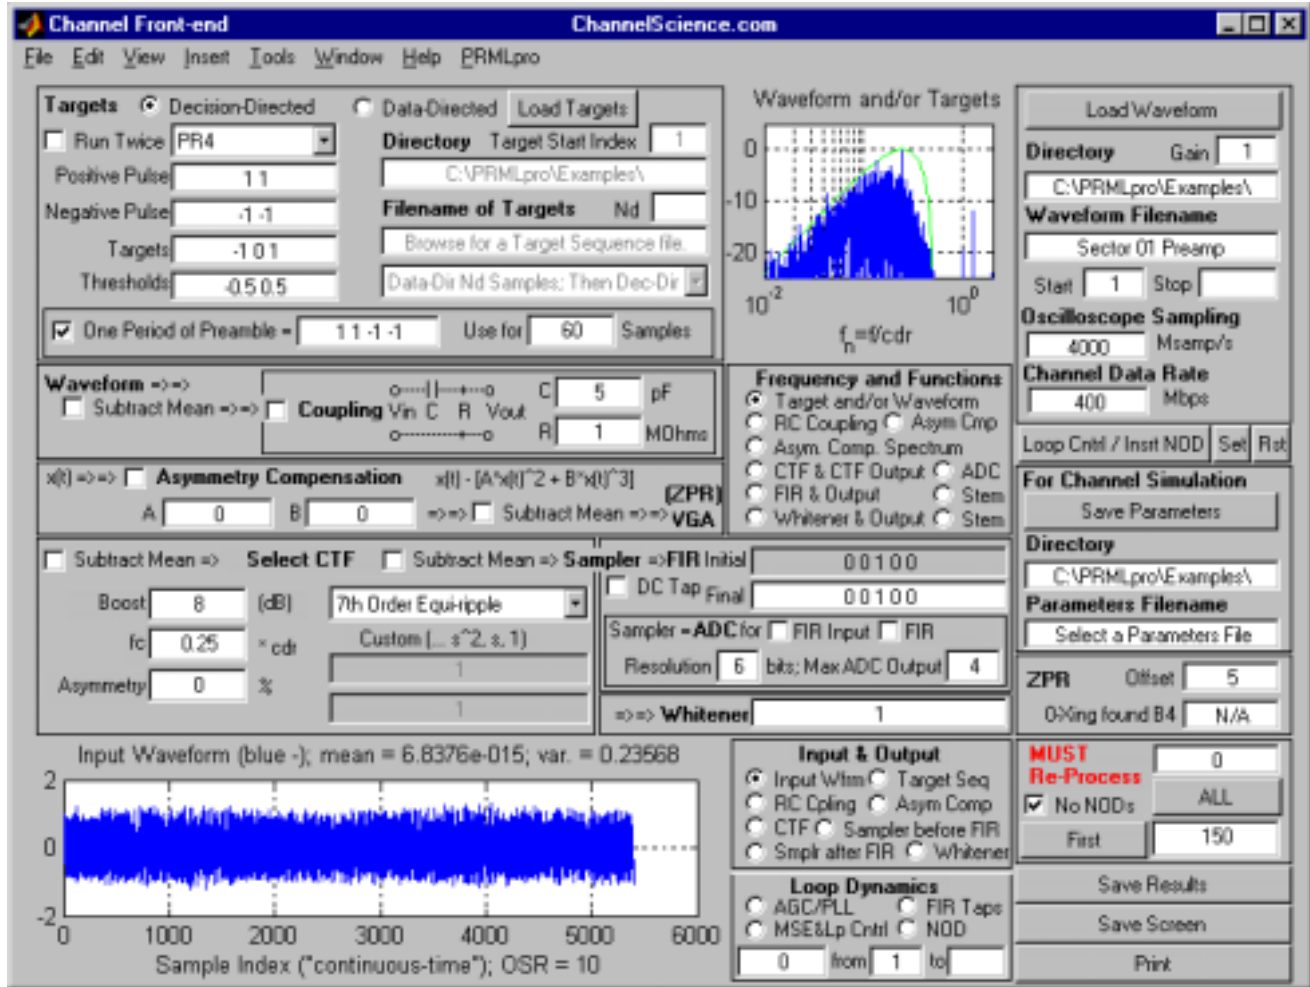


**Figure 2a.** The same waveform from Figure 1 is displayed. The white histogram is calculated using the sample values delimited by the orange vertical markers labeled "Ax" and "Bx." Notice the summary statistics that are automatically provided by the Infiniium at the bottom of the display.

Saved: 30 OCT 2001 15:16:25

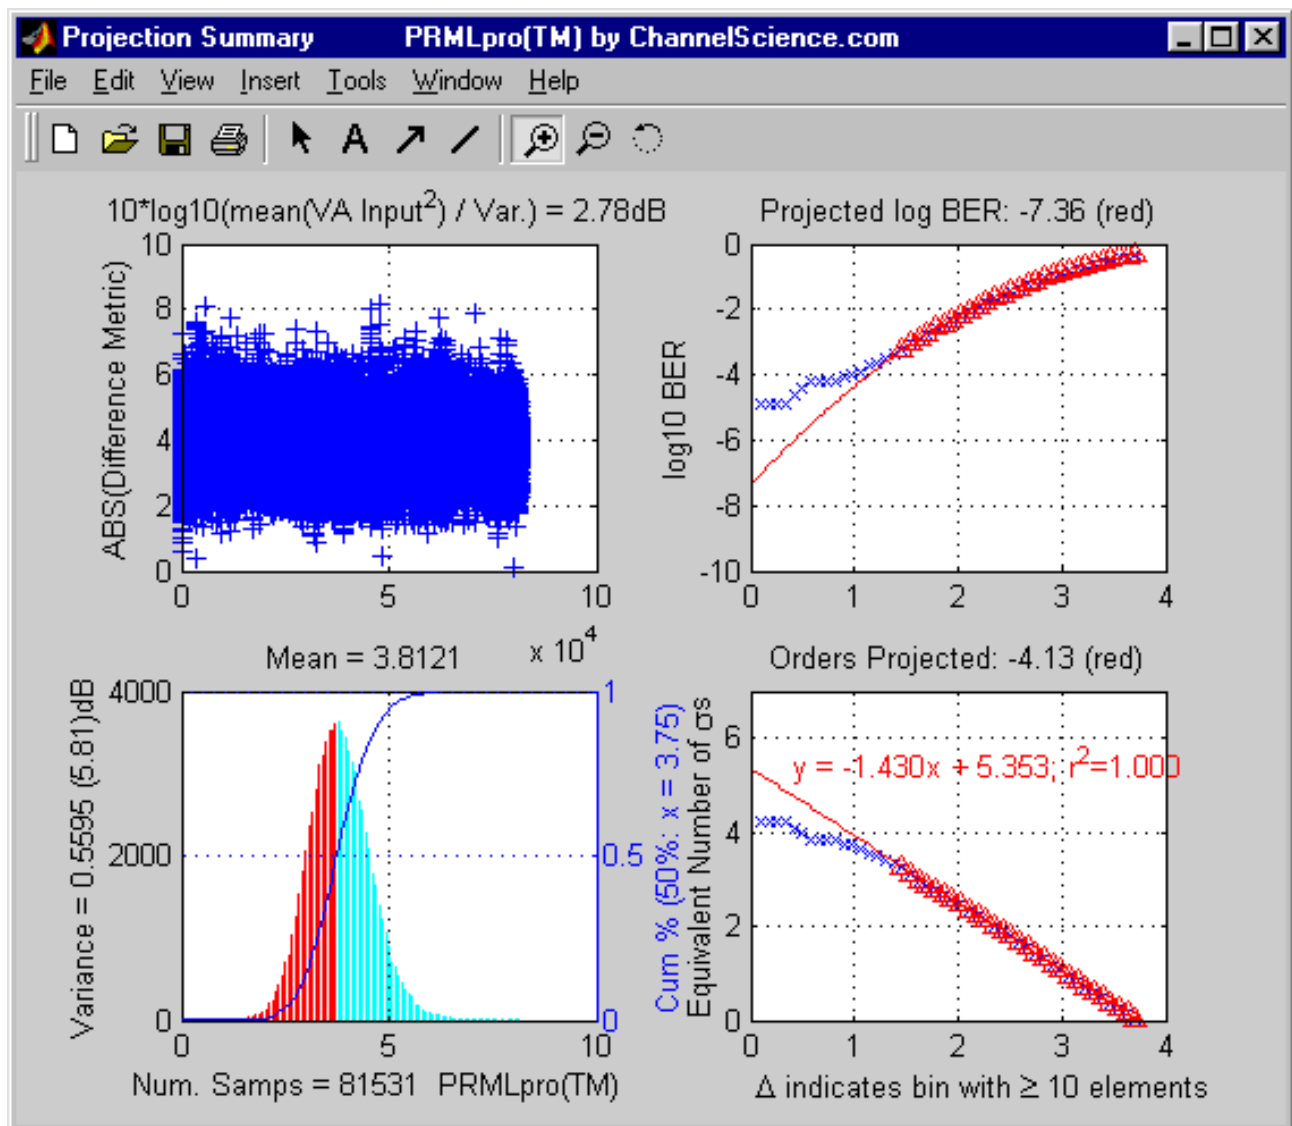


**Figure 2b.** This is the same signal as in Figure 2a, except it is captured *after* the continuous-time filter. Notice that five more-distinct peaks are emerging from the histogram. These five peaks reflect the signal characteristics for which the Viterbi detector was designed.



**Figure 3.** This is the Channel Front-end screen from PRMLpro™, read channel simulation and analysis software from ChannelScience.com. The blue plots are the time-domain and frequency-domain representations of the captured waveform.





**Figure 4.** Upper-left: Scatter plot of the absolute value of Viterbi difference metric values. Lower-left: Their distribution and cumulative sum. Lower-right: Linear fit to the number of  $\sigma$ s from the mean to which each x value corresponds. Upper-right: Calculated BER and projection, based on the linear fit.